

REMARKS/ARGUMENTS

In view of the following discussion, the Applicants submit that none of the claims now pending in the application are anticipated under the provisions of 35 USC § 102(b). Thus, the Applicants believe that all of these claims are now in allowable form.

If, however, the Examiner believes that there are any unresolved issues in any of the claims now pending in the application, Applicants respectfully request that the Examiner telephone Ms. Janet M. Skafar, Esq. at telephone number (650) 988-0655 so that appropriate arrangements can be made for resolving such issues as expeditiously as possible.

Status of Claims

Claims 1-20 are pending in this application.

The Rejection of the Claims Under 35 USC § 102(b)

Claims 1-20 are rejected under 35 USC § 102 (b) as being anticipated by the Osborn et al. patent (U.S. Patent No. 6,249,791), hereinafter referred to as Osborn. Applicants respectfully disagree and traverse the rejection.

Applicants respectfully maintain that Osborn does not teach, explicitly or implicitly, each and every element as claimed. Claim 1 recites: A method of collecting statistics in a database management system comprising a plurality of collection objects, at least one collection object of the plurality of collection objects comprising at least one table, comprising: receiving a workload comprising a plurality of database queries; identifying query statistics for the database queries; consolidating the query statistics for

the database queries to provide consolidated statistics; and generating at least one statistics collection task based on the consolidated statistics.

The rejection asserts that Osborn teaches generating at least one statistics collection task based on the consolidated statistics. The rejection asserts that the summary table recommendations in column 9, lines 29-30 of Osborn is a collection of suggestions. The rejection further asserts that these suggestions are in the form of queries so they can be executed. The rejection contends that a summary table recommendation of Osborn is a statistics collection task of the claimed invention.

Applicants respectfully maintain that the summary table recommendation of Osborn is different from the statistics collection task of the claimed invention. Furthermore, even assuming that a summary table recommendation of Osborn is a suggestion that can be executed and that a statistics collection task is a suggestion that can be executed, does not make a summary table recommendation the same as a statistics collection task.

Per Section 2111.01 of the Manual of Patent Examining Procedure (MPEP) Rev. 5 August 2006, the words of a claim must be given their "plain meaning" unless their plain meaning is inconsistent with the specification. Applicants respectfully maintain that interpreting the term "statistics collection task" as broadly encompassing any suggestion ignores the plain meaning of the language - "statistics collection task." Even if a "statistics collection task" is a suggestion, Claim 1 is not broadly claiming a suggestion. Claim 1 does not recite a "suggestion". Claim 1 is narrowly tailored to recite a "statistics collection task." The plain meaning of a statistics collection task is a task to collect statistics, and this plain meaning is not inconsistent with the specification.

Moreover, even assuming that a statistics collection task and a summary table recommendation are both suggestions, a suggestion for a statistics collection task is

completely different from a suggestion to create a summary table. A statistics collection task and a summary table creation recommendation have completely different purposes with completely different results.

A summary table contains selected subsets of data items, and/or aggregations of data items from one or more source tables of a database. Thus a summary table is a derived table. A particular advantage of summary tables is that they contain pre-generated values, which allow the DBMS to access the summary table to more quickly access and retrieve queried data, rather than accessing the source table(s) of the database.

In contrast, a statistics collection task collects database statistics which are used by the query optimizer to evaluate an execution plan. The statistics collection task gathers meta-data regarding the content of database. In various embodiments, for example, the statistics used by the query optimizer comprise meta-data about the individual tables, columns and column groups within the tables, and indexes on tables within the queries.

A database management system typically provides a tool, for example, RUNSTATS, to collect the database statistics. Thus in various embodiments, a statistics collection task may be a RUNSTATS command.

Further illustrating the difference between a statistics collection task and a summary table recommendation is that statistics to be used by the optimizer to evaluate execution plans may also be collected for a summary table. The optimizer uses the statistics to differentiate execution paths for processing a query, such as a query of data in a summary table.

For the foregoing reasons, Applicants respectfully maintain that Osborn does not teach, explicitly or implicitly, all the recitations of Claim 1. Therefore, Applicants respectfully maintain that Claim 1 is not anticipated by Osborn and is patentable.

Claims 2-8 depend, either directly or indirectly from Claim 1, and are patentable for the same reasons as Claim 1.

Claims 9 and 17 have similar distinguishing recitations as Claim 1. Therefore, Applicants respectfully maintain the Claims 9 and 17 are patentable for the same reasons as Claim 1.

Claims 10-16 depend, either directly or indirectly, from Claim 9 and are patentable for the same reasons as Claims 9. Claims 18-20 depend from Claim 17, and are patentable for the same reasons as Claim 18.

Claims 2, 10 and 18

Claim 2 has the recitation of: wherein the at least one **statistics collection task is to collect database statistics** associated with the at least one table of at the least one collection object of the plurality of collection objects. Applicants respectfully maintain that Osborn has no such teaching.

The rejection cited the following as teaching the recitations of Claim 2. Col. 2 lines 47-50 of Osborn teaches that a database applications layer of a user station of a computer system may be configured to automatically monitor a database and automatically create or delete summary tables from a database. Col. 2, lines 65-67 teaches automatically creating or deleting summary tables based on selected summary table recommendations. Col. 1, lines 29-32 of Osborn teaches that it is not uncommon

for a **database** to include one or more summary tables, which contain selected subsets of data items, and/or aggregations of data items from one or more tables.

Claim 2 recites that the at least one statistics collection task is to collect **database statistics**. As discussed above, the query optimizer uses a set of database statistics to evaluate an execution plan. Various examples of database statistics are shown in the article titled "Tuning SQL with Statistics", by Patrick Bossman, published in the IDUG Solutions Journal, August 1999- Volume 6, Number 2, which is submitted on an Information Disclosure Statement filed on the same date as the instant amendment.

In contrast, and as discussed above, Osborn does not teach the database statistics of the claimed invention. The cited language of Osborn does not teach that the summary table contains data items and/or aggregations of data items that are used by the query optimizer to evaluate an execution plan. The data items and/or aggregations of data items of the summary table of Osborn are the queried data, not statistics or meta-data that is used to evaluate an execution plan.

For the foregoing additional reasons, Applicants respectfully maintain that Claim 2 is not anticipated by Osborn. Claims 10 and 18 have similar recitations as Claim 2 and are patentable for the same reasons as Claim 2.

Claims 5 and 13

Claim 5 has the recitation of: wherein a particular collection object of the plurality of collection objects comprises a plurality of tables, wherein said generating generates a particular statistics collection task to collect database statistics for the plurality of tables of the particular collection object.

The rejection cites column 9, lines 25-33 of Osborn as teaching the recitations of Claim 5. Column 9, lines 25-33 of Osborn teaches:

Based on DBA-supplied criteria and the query records 101 currently residing in the system stat table 100, the respective database applications layer 36 of the user station 24 on which the DBA is operating generates 124 and presents 126 summary table creation recommendations to the DBA. The summary table creation recommendations are preferably returned to the DBA 126 in the form of respective queries the recommended summary tables would be responsive to.

Osborn, column 10, lines 45-65 teaches creating a summary table as follows:

In a preferred embodiment each of the SQL queries stored in a record 101 in the system stat table 100 that is within the selected DBA-defined bounds for summary table creation recommendations, are ranked into the first, second and third tiers according to its corresponding value derived by the above-disclosed algorithm. The respective queries are then presented to the DBA, with the highest ranked query being displayed first, (i.e., at the top of the summary table creation recommendation list supplied to the DBA), and the lowest ranked query being displayed last. In accordance with a further aspect of the present inventions, the respective database applications layer 36 automatically requests the DBMS 34 to create it. In particular, FIG. 6 depicts a method for creating summary tables. To create a new summary table, the respective database application layer 35 generates an SQL query 140 from the query statistics of the record 101 used to generate the particular summary table creation recommendation. This query is then submitted 141 to the DBMS 34 for execution.

Therefore Osborne teaches that a distinct summary table is created based on a distinct SQL statement or query. For example, consider the following two SQL statements:

(1) SELECT T1.C1, T1.C2, T1.C3 FROM T1, T2 WHERE T1.C1=T2.C1
AND T1.Cx=?

(2) SELECT T2.C1, T2.C2, T2.C3 FROM T1, T2 WHERE T1.C1=T2.C1
AND T1.Cx=? AND T1.Cy=?

Using the exemplary SQL statements above, Osborn would recommend creation of two summary tables - one for each SQL statement, even if Tables T1 and T2 are in the same collection object. Osborn does not teach attempting to create one summary table to satisfy both of these SQL statements. Osborn teaches that a separate summary table is created for a distinct query.

In contrast, the present invention of Claim 5 generates a particular statistics collection task to collect database statistics for the plurality of tables of the particular collection object. For example, assuming that Tables T1 and T2 of the two SQL statements above are in the same collection object, the invention would generate a single particular statistics collection task for tables T1 and T2.

For the foregoing additional reasons, Applicants respectfully maintain that Claim 5 is patentable. Claim 13 has similar recitations as Claim 5, and Applicants respectfully maintain that Claim 13 is patentable for the same reasons as Claim 5.

Claims 6 and 14

The rejection asserts that the estimated system cost 114 of executing the SQL query of Fig. 4 of Osborne teaches that each table of the plurality of tables is associated with a table score based on, at least in part, a number of references to said each

table by the queries of the workload (column 9, lines 65-66, the frequency of which respective SQL queries are repeated). Applicants respectfully disagree.

Osborn teaches the frequency with which the same user query has been submitted to the computer system. Osborn does not teach determining a number of references to a table of a query. Osborn does not look into a query to identify the number of references to a table within a query, and therefore does provide a table score of the claimed invention.

For the foregoing additional reasons, Applicants respectfully maintain that Claim 6 is patentable. Claim 14 has similar recitations as Claim 6, and Applicants respectfully maintain that Claim 14 is patentable for the same reasons as Claim 6.

Claims 8 and 16

The rejection asserts that Osborn in column 12, lines 45, 48 teaches the recitation of "when a column specified in the column information is part of an index, generating an index portion of the at least one statistics collection task to collect database statistics for that column. Osborn teaches in column 12, lines 45-49:

generating a summary table creation recommendation based on one or more of the generated statistics for a query submitted to the database system; and creating a summary table based on the summary table creation recommendation.

Osborne does not mention an index. Osborn does not determine if a column of a query is part of an index. Osborn does not generate an index portion of a summary table creation recommendation. Osborn does not disclose the recitation of: when a column specified in the column information is part of an index, generating an

index portion of the at least one statistics collection task to collect database statistics for that column.

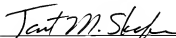
For the foregoing additional reasons, Applicants respectfully maintain that Claim 8 is patentable. Claim 16 has similar recitations as Claim 8, and Applicants respectfully maintain that Claim 16 is patentable for the same reasons as Claim 8.

Conclusion

Consequently, the Applicants believe that all these claims are presently in condition for allowance. Accordingly, both reconsideration of this application and its swift passage to issue are earnestly solicited.

Respectfully submitted,

May 7, 2007



Janet M. Skafar, Attorney
Reg. No. 41,315
Correspondence Customer No. 55070
Telephone: (650) 988-0655
Facsimile: (408) 463-4827